

REPORT DOCUMENTATION PAGE*Form Approved*
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.**1. REPORT DATE (DD-MM-YYYY)**

NOVEMBER 2012

2. REPORT TYPE

CONFERENCE PAPER (PREPRINT)

3. DATES COVERED (From - To)

OCT 2011 – SEP 2012

4. TITLE AND SUBTITLEHardware-based Artificial Neural Networks for Size, Weight, and Power
Constrained Platforms (PREPRINT)**5a. CONTRACT NUMBER**

IN-HOUSE

5b. GRANT NUMBER

61102F

5c. PROGRAM ELEMENT NUMBER

N/A

6. AUTHOR(S)

Bryant Wysocki, Nathan McDonald, and Clare Thiem

5d. PROJECT NUMBER

T1NM

5e. TASK NUMBER

IN

5f. WORK UNIT NUMBER

HO

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)Air Force Research Laboratory/Information Directorate
Rome Research Site/RITB
525 Brooks Road
Rome NY 13441-4505**8. PERFORMING ORGANIZATION
REPORT NUMBER**

N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)Air Force Research Laboratory/Information Directorate
Rome Research Site/RITB
525 Brooks Road
Rome NY 13441-4505**10. SPONSOR/MONITOR'S ACRONYM(S)**

N/A

**11. SPONSORING/MONITORING
AGENCY REPORT NUMBER**
AFRL-RI-RS-TP-2012-063**12. DISTRIBUTION AVAILABILITY STATEMENT**APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA Case Number: 88ABW-2012-5205
DATE CLEARED: 01 OCT 2012**13. SUPPLEMENTARY NOTES**

This paper has been submitted to IEEE Symposium on Computational Intelligence for Security and Defense Applications (CISDA), Singapore, 15-19 April 2013. This is a work of the U.S. Government and is not subject to copyright protection in the United States.

14. ABSTRACT

A fully parallel, silicon-based artificial neural network (CogniMem CM1K) built on zero instruction set computer technology was used for change detection and object identification in video data. Fundamental pattern recognition capabilities were demonstrated with reduced neuron numbers utilizing only a few, or in some cases one, neuron per category. This simplified approach was used to validate the utility of few neuron networks for use in applications that necessitate severe size, weight, and power restrictions. The limited resource requirements and massively parallel nature of hardware-based artificial neural networks make them superior to many software approaches in resource limited systems, such as micro-UAVs, mobile sensor platforms, and pocket-sized robots.

15. SUBJECT TERMS

Artificial neural networks; size, weight and power (SWaP); radial basis function; zero instruction set computing (ZISC); pattern recognition

16. SECURITY CLASSIFICATION OF:**17. LIMITATION OF
ABSTRACT****18. NUMBER
OF PAGES****19a. NAME OF RESPONSIBLE PERSON**

CLARE D. THIEM

a. REPORT

U

b. ABSTRACT

U

c. THIS PAGE

U

UU

7

19b. TELEPHONE NUMBER (Include area code)

N/A

Hardware-based Artificial Neural Networks for Size, Weight, and Power Constrained Platforms

B. T. Wysocki, N. R. McDonald, and C. D. Thiem
Air Force Research Laboratory, Information Directorate Rome, NY

Abstract— A fully parallel, silicon-based artificial neural network (CogniMem CM1K) built on zero instruction set computer technology was used for change detection and object identification in video data. Fundamental pattern recognition capabilities were demonstrated with reduced neuron numbers utilizing only a few, or in some cases one, neuron per category. This simplified approach was used to validate the utility of few neuron networks for use in applications that necessitate severe size, weight, and power restrictions. The limited resource requirements and massively parallel nature of hardware-based artificial neural networks make them superior to many software approaches in resource limited systems, such as micro-UAVs, mobile sensor platforms, and pocket-sized robots.

Keywords -- *artificial neural networks; Size, Weight and Power (SWaP); radial basis function; zero instruction set computing (ZISC); pattern recognition*

I. INTRODUCTION

There exists a potential for disruptive enhancements to Department of Defense (DoD) capabilities and efficiencies through the development of autonomous human-system interactions. Intelligent platforms that collaborate with their human operators have the capacity to enhance and complement the human capability while reducing tedium and adding resiliency and adaptability to current systems. Programmable machines are limited in their ability to address such fuzzy combinatorially complex scenarios. Neuromorphic processors, which are based on the highly parallelized computing architecture of the mammalian brain, show great promise in providing the environmental perception and comprehension required for true adaptability and autonomy.

The increasing resolution and speed of today's advanced sensor platforms provide an overwhelming and exponentially growing supply of data, which has subsequently created a demand for autonomous pattern recognition systems that scour raw or preprocessed data in an effort to extract meaningful information [1 - 3]. The improved processing power of modern high performance computers enables implementation of large, sophisticated pattern recognition systems based on statistical analysis and neural network schemes. This latter approach is the subject of several large-scale efforts to create electronic systems that mimic behaviors in the brain. This fundamentally different approach, frequently referred to as neuromorphic computing, is thought to be better able to solve fuzzy perception and classification problems historically difficult for traditional, von Neumann-based computers.

The neuromorphic community was revitalized when, in 2008, memristive devices were brought to public attention by

Hewlett Packard (HP) [4], though devices possessing similar behavior, as predicted by Leon Chua in 1971 [5], had been observed since the 1970s [6 - 9]. These "memory resistors" possesses several attributes which make them effective hardware incarnations of biological synapses. First, these two-terminal devices function as variable resistors, satisfying the "learning" requirement seen in real synapses. Additionally, these devices are non-volatile. Unlike transistors, each memristive device's state persists even in the absence of applied power. Lastly, these devices exhibit this behavior even in the nanoscale regime. For traditional computing, these properties could lead to instant-ON computers or non-volatile field-programmable gate arrays (FPGAs). These devices potentially offer to neuromorphic circuit engineers the means to achieve the density, reconfigurability, and low power requirements needed to build analog neural circuitry. Memristive technology, however, must mature before it can be utilized in such designs. Memristor-built non-volatile memory is expected to be commercially available by 2014 [10], and HP predicts that memristive memory will eventually replace FLASH, solid state, and DRAM as a universal memory format [11]. Due to the disruptive nature of technology developments and to the exponential growth of technology as a whole, it is difficult to predict when, or if, memristive technology will enable neuromorphic thinking machines. Recent advances in silicon-based artificial neural networks (ANNs) offer an alternative approach with many of the capabilities desired in future memristive systems but which are available now, CMOS compatible, and inexpensive.

Neuromorphic computing goals such as emulating mammalian brains prove daunting due to the processing power required to emulate all 4 million neurons of even a mouse's brain, much less that of a household cat, with approximately 300 million neurons. At the same time, extraordinary examples of pattern recognition and behavior are evident throughout the animal kingdom with significantly fewer neurons. For example, the roundworm, with 302 neurons and 8,000 synapses, can sense and track waterborne chemical signatures and navigate towards their locations [12]. We have found in our experimentation with hardware-based neural networks that useful applications can be realized with relatively few active neurons. In fact, in one instance, only a single neuron was required to enable relevant change detection in a video surveillance system. Limited neuron approaches employed near the sensor may be used to reduce large data sets, saving critical transition bandwidth while reducing the burden on analysts and system operators.

Hardware-based ANNs are ideally suited for mobile or portable platforms with strictly limited size, weight, and power

(SWaP) resources. The robotics industry (currently a \$10 billion enterprise) is poised for exponential growth with an expected commercial market of over \$15 billion by 2015 [13]. As an enabling technology, ANNs will play a substantial role in the development of autonomous and semi-autonomous robots for use in industrial, commercial, and military markets. The demand for small, mobile, battery-powered systems will favor low-energy hybrid processing units consisting of both standard microcontrollers and hardware-based neural networks on the same chip. The parallel nature of these neural co-processors makes them superior over software techniques for SWaP restricted applications.

Our research focuses on the implementation of real-time ANN pattern recognition in platforms with severe SWaP constraints, such as micro-UAVs, mobile sensor platforms, and pocket-sized robots. These restrictions practically rule out traditional software approaches which often run too slowly due to the inherent serial nature of von Neumann architectures or require high performance processing for operation. Hardware implementation provides a reduced footprint with the additional benefit of massively parallel execution. While nano-enabled neuromorphic architectures are extremely promising, their realization will take time. Meanwhile, there exist commercially available technologies today that offer partial solutions. In particular, parallel processing capabilities are afforded by FPGAs [14] and general-purpose computing on graphics processing units (GPGPUs). With the recent availability of application-specific integrated circuits (ASIC) based on zero instruction set computing (ZISC), not only is there an even greater reduction in footprint and power but also native support for massively parallel operation. It is technically feasible using current state-of-the-art fabrication techniques at the 22 nm node to manufacture ASIC ANN chips approaching 500,000 parallel neurons.

This paper will first examine the current state of memristive development with emphasis on architectural and fabrication challenges followed by a review of the technical aspects of the CM1K ASIC chip and its ZISC operation. Next, two experiments demonstrating change detection and object recognition in live video feed will be presented. Lastly, the results of these experiments will be used to weigh in on two competing perception paradigms: few sensors/complex computations, and many sensors/simple computation.

II. CHALLENGES WITH NANO-ENABLED NEUROMORPHIC CHIPS

A wide variety of technologies can be classified as “memristive devices” including resistive random access memory (ReRAM), phase change RAM (PCRAM), magnetoresistive RAM (MRAM), and spin-transfer torque MRAM (SST-RAM). While all these devices operate under different physical principles, they all possess two key attributes: 1) variable resistance and 2) non-volatility. Despite all the progress that has been made in memristive devices over the past several years, commercial memristive device products are still unavailable.

There are two critical tasks for successful memristive device integration with CMOS: manufacturability and usability. Concerning the former, the devices to be used must consist of materials that are permitted inside a CMOS foundry, which further restricts the materials allowed in the front end of line (FEOL) as compared to the back end of line (BEOL). All the processing steps needed to make the devices’ structure must be scalable to fabricate devices en masse. Lastly, the devices must be all functionally identical (though some applications may actually exploit device non-uniformities). Part of the difficulty of manufacturing memristive devices is that the physics of device switching is not well understood at nanometer size scales. In particular, ReRAM (of which PCRAM is a subset) may be composed of binary metal oxides, chalcogenides, or perovskites, among other materials, and switch due to filament formation, vacancy migration, phase change, or other processes [15]. At these scales, small variations in the device size or material composition often have large effects upon subsequent device switching parameters.

However, because of this variety of materials and mechanisms, different device resistance values, switching voltages, and switching times are available to the circuit designer. When considering the appropriate device metrics of reliability and endurance that must be attained, one must first consider the intended use of the device. For von Neumann computing applications, if these devices are to replace Flash or SRAM, then endurance cycles of about 10^6 and write speeds of a couple tens of nanoseconds must be achieved, respectively. Even if memristive devices cannot meet these requirements, SWaP savings may still be achieved by strategically replacing some transistors in a circuit. For devices used in neuromorphic applications, the range of addressable resistance values and the operative voltages will be more critical than the write speed. Because of these varied ends, there will likely be a variety of memristive device “flavors” available to the circuit designer in the future.

In the meantime, ZISC neural networks may be usable to partially replicate some of the anticipated advantages of a memristive device-based hardware neural network. In this way, learning algorithms and neural network structures may be further developed now and subsequently mapped onto memristive device technology if and when it becomes available.

III. A HARDWARE-BASED ARTIFICIAL NEURAL NETWORK

A. Chip Specifications

A fully parallel, silicon-based neural network chip (CM1K) developed by CogniMem Technologies Inc., based on IBM’s earlier series of ZISC chips [16,17], was used in this work. The CM1K is configured with two available types of non-linear classifiers: a Radial Basis Function Network (RBF) and a K-Nearest Neighbor classifier (KNN). The chip possesses 1024 neurons, each with its own memory for trained signature

storage and a processor for recognition and distance calculations. The memory within every neuron contains 256 elements, each with an 8-bit capacity for a total of 256 bytes of information per neuron. The identical neurons learn and respond to vector inputs in parallel while they incorporate information from all the trained neurons in the network through a bi-directional parallel neuron bus. Execution of the recognition logic is independent of the number of participating neurons, and multiple chips can be cascaded in parallel for scalable implementation. Figure 1 shows the general topology of such a restricted coulomb energy network. CogniMem recently demonstrated a cascaded network of 100 chips with over 100,000 parallel neurons, all contained within 1/10 of a cubic foot and consuming less than 20 Watts of power yet performing at a level equivalent to 13.1 Teraops of pattern recognition performance [18]. Additional details regarding CM1K operation and architecture may be found in [19,20]




Figure 1: Neural Network Diagram. Each input node accepts a maximum of 256 elements (x_N), each with 8-bit resolution. These are fed in parallel to up to 1024 neurons. All recognition events are passed through to the output layer with the associated category and confidence level.

B. Operational Parameters

In such an architecture, the operational status of each neuron can be in one of three possible states: idle, ready-to-learn, and committed. The idle neurons are empty of knowledge but can be trained sequentially with the next neuron in the chain configured in the ready-to-learn state. Once a neuron is trained it becomes committed and any pre-existing influence fields are adjusted to accommodate the new knowledge. During recognition, the input vector is passed to all the committed neurons in parallel, where it is compared to the stored vector or trained prototype. If the distance between the input vector and a neuron prototype falls within the influence field, the neuron “fires” generating local output signals consisting of fire flag, signal distance, and category type. In the case that no neurons fire, the input signal can be used to train the ready-to-learn neuron with the unrecognized signature. This provides the means for recognition and training to be accomplished simultaneously. A sample visual cue and grey scale image of the vector signature within a neuron’s memory is given in Figure 2.

IV. PLATFORM DESIGN

The video surveillance system is designed to monitor a live video scene and alert an analyst when an event of interest (an intrusion) has occurred. Our system utilizes the CM1K chip as part of the CogniMem V1KU evaluation module specifically designed for video and image prototype development. Our control software communicates with the V1KU through a USB connection to a microcontroller. A graphical user interface (GUI) was built in JAVA to facilitate training and evaluation. The system does not transmit or store the active video feed until the ANN has determined that an intrusion event has occurred with the intent to optimize the transmission bandwidth for mobile or wireless platforms.




Figure (2): a) A captured video frame showing the subject in its input resolution. b) A pictorial representation of the neuron content trained with the above image.

The program scans at 60 frames per second live video feed and compares the region of interest to those trained into the neurons memory. Fine tuning of neuron sensitivity for a specific signature can be manually adjusted by adjusting the neuron active influence field or the distance from ideal, where a neuron will still recognize the target as a specific category. The distances can be calculated using one of two norms: the Manhattan method or the L_{sup} method. In the Manhattan method,

$$D_{Man} = \sum_{i=1}^n |V_i - P_i|, \quad (1)$$

where D_{Man} is the sum of the differences between n dimensional vector signatures V_i and P_i . In the L_{sup} method,

$$D_{Lsup} = \text{Max}|V_i - P_i|, \quad (2)$$

where D_{Lsup} is the maximum separation of V_i and P_i .

Manhattan distances were chosen for our design to emphasize the general differences between signatures with equal weights on all components. A neuron fires when the input vector lies within a specified distance, that is, falls within the influence field of a neuron in the decision space.

V. EXPERIMENTS IN MINIMAL NEURON REQUIREMENTS

A. Changes Detection in Live Video

We examined the problem of change detection in video surveillance with the intent of utilizing the least amount of processor resources. We started with the straightforward task of monitoring the entry point into a room for any activity. A single neuron was trained to recognize an image of the entry point. The neuron's influence field was manually adjusted to the sensitivity required to detect a subtle change to the field of view. The native monochromatic video feed to the CM1K was progressive scan at 752x480 pixel resolution and 60 frames per second while the neuron memory was scaled down to 187 elements each with 8-bit depth. While maintaining the context of the broadcast vector, it required $N+2$ clock cycles to pass the input to the network. With the CM1K clocked at 27 MHz or 37 ns per clock cycle, the broadcast of our 187 element input vector took only 7.0 μ s. A block diagram of the complete system configuration is shown in Figure 3.

Despite the low resolution of the stored prototype (187 bytes), the system reliably responded to changes in the camera's field of view (FOV) or to more localized changes in a region of interest within the FOV. In this case, the system was programmed to alert a human analyst of any change and to capture images of the disturbance for review. Data capture and transmission continued until the intrusion moves outside the sensor's FOV. It is important to mention that although we implemented our system under the control of a personal computer, the ZISC chip learns and recalls patterns without internal code or the need for constant external supervision.

This simple implementation of a single neuron for video change detection proved to be very reliable over a 36 hour period of entryway monitoring identifying 26 events with zero false positives and zero missed occurrences.

B. Specific Target Recognition

Pattern recognition in complex scenes often plagues ANNs, since a subject's spatial orientation along with environmental variables such as lighting and background affect the system's ability to accurately perceive the target. These inconsistencies can be addressed using three basic techniques: increasing the number of neurons to account for variability, preprocessing the video to reduce variance, and controlling the environment or setting of the scan. While applications exist where control over situational effects can be adequately controlled by engineering the platform's environment [20], reducing the variance in video

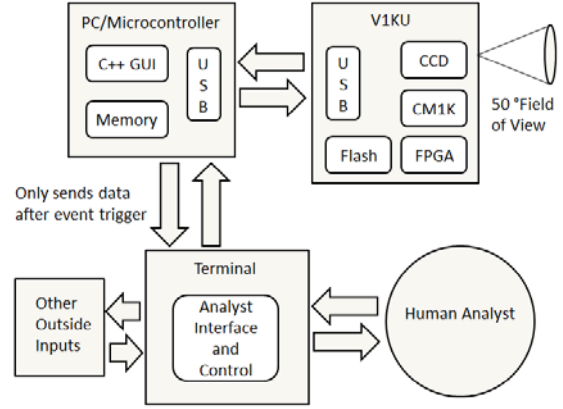


Figure 3: A block diagram depicting the system configuration.

streams can be particularly challenging due to the inherent inclusion of erroneous and unpredictable background effects [21].

In this portion of the experiment, we tested the system's ability to detect specific targets in a controlled setting. Simulated vehicle traffic was monitored at the entrance point of a scaled model parking facility. A video camera monitored the incoming traffic (remote controlled cars), and the system alerted an attendant when either an unrecognized vehicle or a prespecified vehicle approached the gate. For this experiment, a pool of four vehicles under constant lighting was used. A single neuron was sufficient to identify a particular vehicle (Figure 4). Since these neurons do not interpolate data, all flagged images had to be consistent with the orientation of the trained image.

With as few as one neuron per vehicle the system accurately identified each of the four vehicles and subsequently notified the analyst when a specifically flagged or unrecognized vehicle approached. Additional training was not required as long as the environmental aspects were held constant but improved reliability when environmental controls were lessened.

The number of neurons required to distinguish N objects did not grow linearly but at a much faster rate, such that 257 neurons were required to properly categorize 34 distinct targets under very controlled conditions. The realm of applications within a single processor's 1024 neuron capacity is significant but the need for additional performance through cascaded chips is required for complex relationships. Using this platform, we plan to construct a 100,000+ parallel neuron network for additional research.

VI. RESULTS AND DISCUSSION

These two simplified tests illustrate that useful ANN systems can be designed to operate with extremely limited resources for use in SWAP constrained platforms. A nonlinear relationship was found between the number of prototype categories and the number of neurons required for

identification, highlighting the need for increased density, low-power chip designs.

SRAM currently serves as the neuron's memory in this integrated circuit which consumes a significant amount of wafer real-estate and requires substantial energy resources for periodic refreshing. SRAM typically uses six transistors for storage and control of a single bit of CMOS memory. Additionally, since all the neuron knowledge is presently stored in SRAM, it must be saved off-chip in Flash when power is removed. Such requirements severely reduce the achievable density in SRAM-based neural networks. As mentioned previously, memristive devices hold great potential for the development of ultra-high density memory with reduced resource requirements. The non-volatile nature of memristive memory could lead to super-efficient ANNs that lay dormant without consuming power, ready to instantly respond when needed. By selectively replacing certain SRAM and DRAM transistors with CMOS-compatible memristive devices, ANNs can subsequently achieve increased density, reduced power, and instant-ON capabilities.

As was shown in both of the previous examples, a single neuron in this system was sufficient to achieve reasonable discrimination for pattern recognition. Coarse resolution in trained neurons may not necessarily be something that must be improved upon. In general, modern environmental processing relies heavily upon a small number of high resolution sensors and a computer capable of solving complex differential equations in real time. It is very unlikely that flying insects are performing such computationally intensive sensor processing in their brains. Rather, such biological systems perform simple computations using measurements from numerous crude sensors. This is called the sensor-rich feedback control paradigm. For example, it is thought that a fly determines a global vector representing how the fly is moving with respect to its environment. From observed vector patterns, select neurons fire when preferable flight directions are identified [22]. Extending this many sensors/limited computation paradigm to autonomous systems is a natural progression towards bio-inspired computation.

VII. CONCLUSIONS

A fully parallel, silicon-based ANN was used to monitor video data. In this work, change detection and simple object recognition were demonstrated with reduced neuron numbers utilizing only a few, or in some cases one, neuron per category. This simplified approach was used to validate the utility of few neuron networks for use in applications that necessitate severe SWAP restrictions. The limited resource requirements and massively parallel nature of hardware-based ANNs make them superior to many software approaches in such resource limited systems, such as micro-UAVs, mobile sensor platforms, and pocket-sized robots. These fully CMOS compatible designs will likely play a substantial role in the development of semi-autonomous robotic platforms. Configurations having multitudes of crude sensors connected to layered ANNs will more closely emulate the structure of biological systems and may outperform systems which rely upon brute forcing

complex equations upon data from a few high resolution sensors.

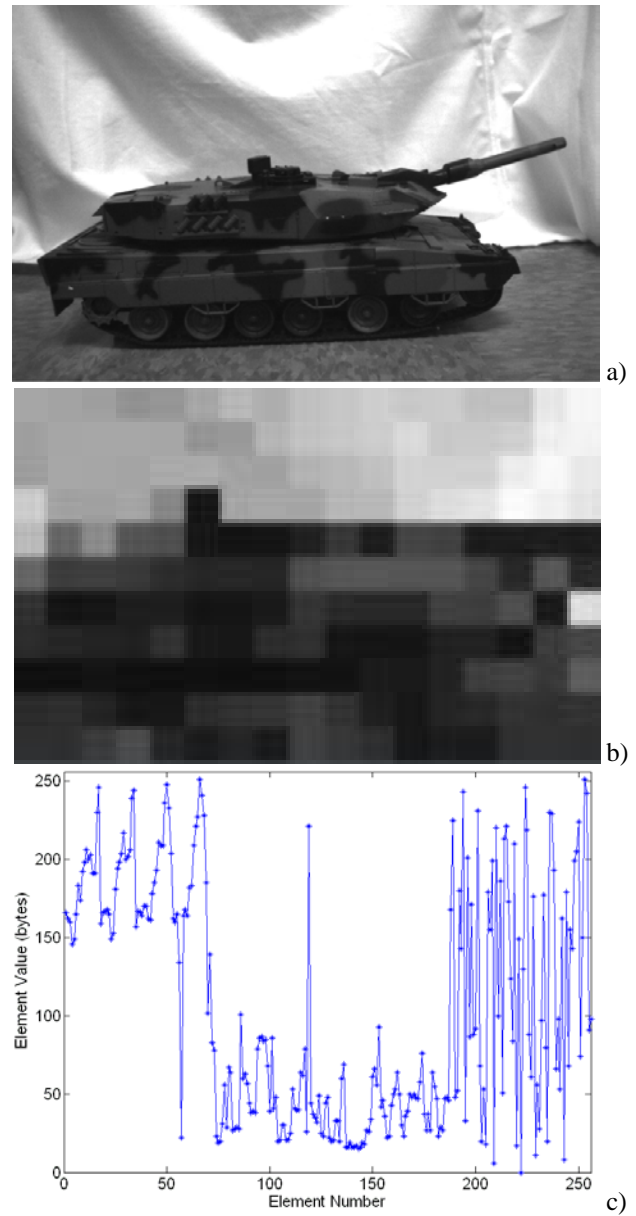


Figure 4: a) Profile view of one of the vehicles recognized by one neuron. b) A pictorial representation of the neuron content trained with the above image. c) A plot of the neuron's prototype vector.

ACKNOWLEDGMENTS

This work was performed at Air Force Research Laboratory Information Directorate, Rome, NY leveraging funding from AFOSR's test and Evaluation Program. Approved for Public Release (PA Case # 88ABW-2012-5205 01 Oct 2012); distribution unlimited.

REFERENCES

- [1] P. W. Singer, *Wired for War: The Robotics Revolution and Conflict in the 21st Century*, Penguin, 2009.
- [2] W. J. A. Dahm, "Report on Technology Horizons: A Vision for Air Force Science & Technology During 2010-2030," [Online], Available: URL: http://www.aviationweek.com/media/pdf/Check6/USAF_Technology_Horizons_report.pdf [Accessed Jan 04, 2012].
- [3] J. Misra and I. Saha, Artificial neural networks in hardware: A survey of two decades of progress, *J. Neurocomput.* vol 74. 2010.
- [4] D.B. Strukov, G.S. Snider, D.R. Stewart, and R.S. Williams, "The missing memristor found," *Nature*, vol. 453, May. 2008, pp. 80-83.
- [5] L. Chua, "Memristor - the missing circuit element," *IEEE Trans. Circuit Theory* 18, 507-519 (1971).
- [6] Y. Hirose and H. Hirose, "Polarity-dependent memory switching and behaviour of Ag dendrite in Ag-photodoped amorphous As₂S₃ films," *J. Appl. Phys.*, Vol. 47, No. 6, p. 2767, 1976.
- [7] A. Beck, J. G. Bednorz, Ch. Gerber, C. Rossel, D. Widmer, "Reproducible switching effect in thin oxide films for memory applications," *Appl. Phys. Lett.*, Vol. 77, No. 1, p. 140, 2000.
- [8] P. van der Sluis, "Non-volatile memory cells based on ZnxCd1-xS ferroelectric Schottky diodes", *Appl. Phys. Lett.*, Vol. 82, No. 23, p. 4089, 2003.
- [9] S. Seo, M. J. Lee, "Reproducible resistance switching in polycrystalline NiO films," *Appl. Phys. Lett.*, Vol. 85, No. 23, p. 5655, 2004.
- [10] C. Mellor, "HP's faster-than-flash memristor at least TWO years away
Plus: Storage boffins discuss photonic chip comms," *The Register* [Online]. Available: http://www.theregister.co.uk/2012/07/09/hp_memristor_and_photons/ [Accessed 19 September 2012].
- [11] J. Nickel, *Memristor Memory: A Fundamental Shift*, 2011 IEDM, 2011.
- [12] K. Shen and C. I. Bargmann, "The immunoglobulin superfamily protein SYG-1 determines the location of specific synapses in *C. Elegans*," *Cell*, v.112, n.5, pp. 619-630, 2003.
- [13] ABI Research. Personal Robots Are Here (and by 2015 They'll Be Worth \$15 Billion). [Online]. Available: <http://www.abiresearch.com/abiprdisplay.jsp?pressid=1023> [Accessed Jan 03, 2012].
- [14] A. R. Omondi and J. C. Rajapakse, *FPGA Implementations of Neural Networks*, Springer, 2006.
- [15] Pershin, Y. V.; Di Ventra, M. *Adv. Phys.*, 60, 145-227, 2011.
- [16] Å. Eide, T. Lindblad, C. S. Lindsey, M. Minerskjöld, G. Sekhniaidze, and G. Székely, "An implementation of the zero instruction set computer (ZISC036) on a PC/ISA-bus card", *WNN/FNN*, , 1994.
- [17] F. M. Dias, A. Antunes, A. Mota, "Artificial Neural Networks: a Review of Commercial Hardware", *Engineering Applications of Artificial Intelligence*, IFAC, Vol. 17/8, pp. 945-952, 2004.
- [18] Cognimem Technologies, Inc. Cognimem Communique Volume 1, Issue 2 [Online], Available: URL: <http://general-vision.com/docs/Newsletters/CogniMem%20Communique%27,%20Vol%201,%20Issue%202.pdf> [Accessed Jan 10, 2012].
- [19] J. Y. Boulet, D. Louis, C. Godefroy, A. Steimle, P. Tannhof, G. Paillet, Patent US5621863 [Online], Available: URL: <http://patft.uspto.gov/netacgi/nph-Parser?Sect2=PTO1&Sect2=HITOFF&p=1&u=/netahtml/PTO/search-bool.html&r=1&f=G&l=50&d=PALL&RefSrch=yes&Query=PN/5621863> [Accessed Jan 10, 2012].
- [20] Y. Q. Liu; D. Wei; N. Zhang; M. Z. Zhao; , "Vehicle-license-plate recognition based on neural networks," *Information and Automation (ICIA)*, 2011 IEEE International Conference on , vol., no., pp.363-366, 6-8 June 2011.
- [21] W. S. Leputra, T. Tan, and S. Venkatesh, *Unified 2D-3D Video Scene Change DetectionFramework for Mobile Camera Platforms*, 11th Int. Conf. Control, Automation, Robotics and Vision, 2010.
- [22] R. Zbikowski, "Fly Like a Fly," [Online]. Available: URL:<http://spectrum.ieee.org/computing/hardware/fly-like-a-fly> .